

100
102
104
200
108

The diagram illustrates a distributed query processing architecture. At the top, three groups of components are labeled: 'Queries' (100), 'Front End (FE) Processors' (102), and 'Back End Processors (P)' (108). Below these, a grid of boxes represents the internal structure. On the left, multiple input arrows point to a row of boxes labeled 'FE(1)', 'FE(2)', '...', 'FE(j)', and 'FE(i)'. From each of these boxes, multiple arrows point to a column of boxes labeled 'QI(1)', 'QI(2)', '...', 'QI(m)'. Finally, from each of these boxes, multiple arrows point to a row of boxes labeled 'P(1)', 'P(2)', 'P(3)', 'P(4)', '...', and 'P(n)'. The connections are fully meshed between all components.

FIG. 1

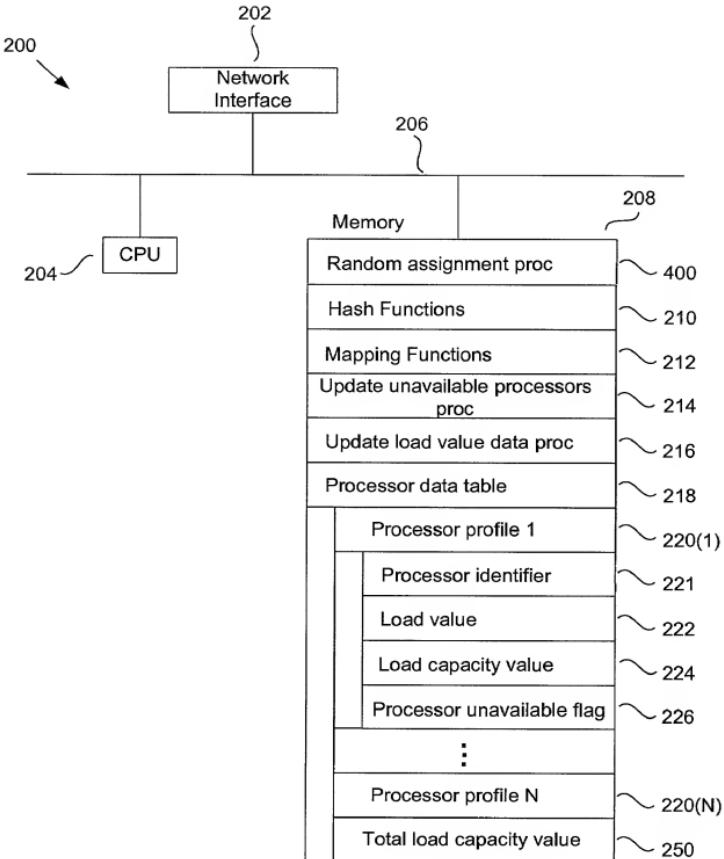


FIG. 2

Processor identifier	Load value	Load capacity value	Unavailable data
1	3	250	
2	2	800	
3	4	350	
4	1	400	
5	6	100	
6	9	300	
⋮			
N			

Total load capacity value: 12, 500

Fig. 3A

Processor identifier	Load value	Load capacity value	Unavailable data
1	3	250	
2	2	800	
3	4	350	
4	1	400	✓
5	6	100	✓
6	9	300	
⋮			
N			

Total load capacity value: 12, 000

FIG. 3B

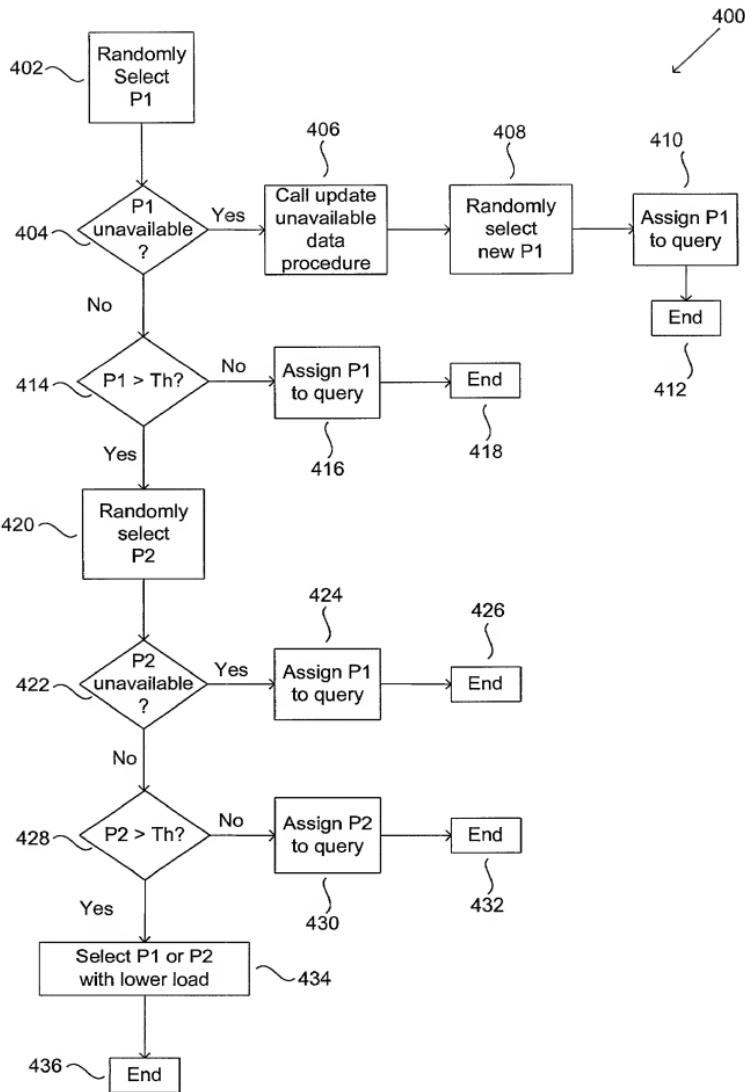


FIG. 4

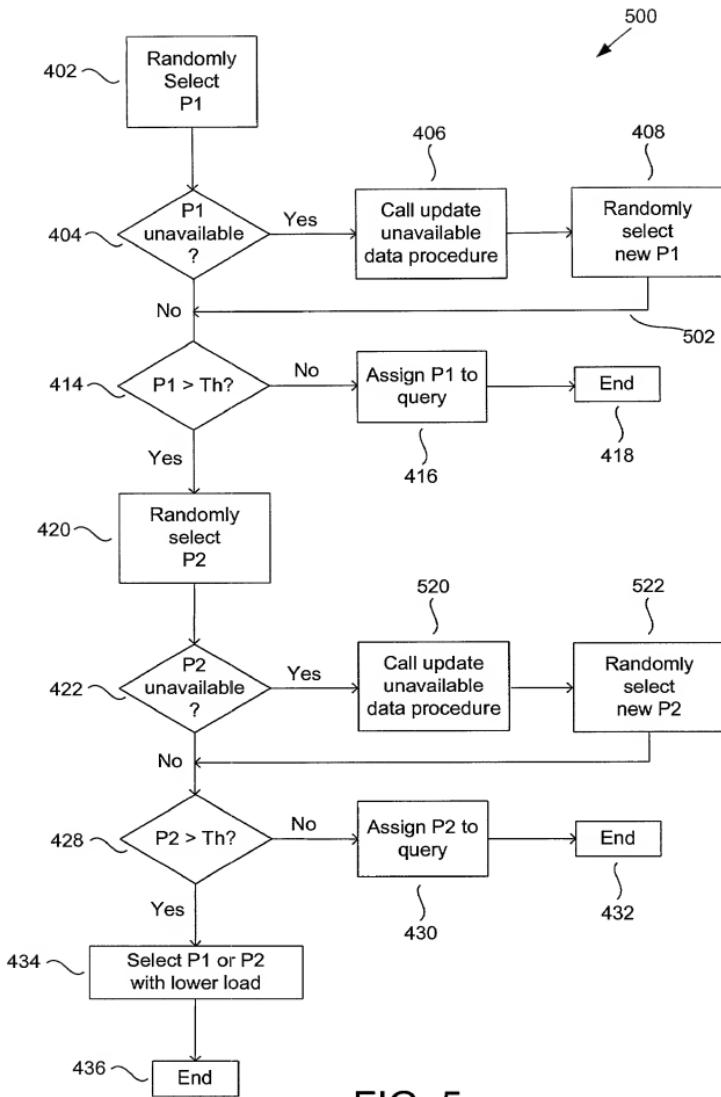


FIG. 5